Intel Developer Update is Intel's monthly online news magazine for developers. As the official publication of developer.intel.com, it brings hardware, software, and Web developers the latest information on Intel initiatives, technologies, platforms, and products.

Cover Story
Each month, we run a cover story on the most significant industry announcement, trend, or development for the month.

Top Stories
The issue's Top Stories are the three or four articles that cover the month's hottest topics. For a list of all the articles in the issue, see the Table of Contents.

Contact the Editor
To make *Intel Developer Update* a better information resource, we invite you to share your thoughts on what we've published or what you'd like to see covered. Comments are always welcome.

Archives
Our archives contain two groups of previously published articles. One group contains all the articles that appeared in *Platform Solutions News*, the earlier version of *Intel Developer Update*. The articles date from September 1997 through August 1999. The other group is set up to contain *Intel Developer Update* articles dating from the inaugural September/October 1999 issue.

Bookmarking
We advise against bookmarking article pages. They're accessible online only during the month the issue is live. Thereafter, they're removed to our archives. Instead, we suggest that you bookmark the PDF (Adobe* Portable Document Format) file versions of the articles. You'll find buttons for the PDF files labeled "print article" in the right navigation section of each article. A PDF for the entire issue is labeled "print magazine" and is located near top right side of the IDU home page.

Table of Contents

(Click on page number to jump to articles)

COVER STORY

Bring Your Team to Fall IDF 3

COLUMNS

From the Editor 6

This Much We Owe to Our Friends 8

DEPARTMENTS

APPLIED COMPUTING

Optimizing Performance in Virtual Private Networks 10

Motherboards: An Answer to Industry Demand 14

INITIATIVES AND TECHNOLOGIES

Zero to OS in 20.2 Seconds 17

Serial ATA: An Evolutionary Transition 20

Intel Enables Universal Plug and Play for Linux* 24

The Intel® Itanium™ Processor and Security 28

Cover Story

Bring Your Team to Fall IDF

David Barkai
IDF Content Planning Team
Intel Architecture Marketing Group
Intel Corporation

Vincent Merrick
IDF Content Planning Team
Intel Architecture Marketing Group
Intel Corporation

Overview

Get ready to power your development program and connect with the latest platform technologies and software solutions at the Intel Developer Forum Conference, Fall 2000 (IDF) happening August 22 through 24 in San Jose, California. With the theme of "Powering the net. Connecting the net." this IDF conference is a must-attend event, and it's almost here.

You should make plans now to be in attendance for a significant announcement on Intel's entry into wireless communications, the latest developments in IA-64 and IA-32, and more software solutions than an IDF conference has ever delivered before.

With the equivalent of 40 days of content packed into a three-day event, IDF Conference Fall 2000 will deliver so much information that your organization will need to send a team to cover it all. IDF is now bigger than ever, with 25 percent more sessions than the last conference, including more than 100 hours devoted to software and 140 hours focused on hardware technologies. This is your opportunity to hear more than 150 industry leaders describe the latest net-centric technologies, products and tools. In addition to technical tracks and hands-on Labs, IDF is expanding the Fall conference to include over 60 percent more hours at the Demo Showcase, compared with the Spring conference.

To make room for the added content, IDF is moving to a new venue, the San Jose Convention Center. For the first time, software developers will have their own dedicated site, the beautiful Fairmont Hotel in downtown San Jose.

Intel CEO Craig Barrett kicks off the conference with an opening keynote that shares Intel's vision of platform leadership, communications, and enabling solutions to power and connect the net. Other keynote speakers include: Albert Yu of the Intel Architecture Group on the vision for client and server platforms; Pat Gelsinger, Chief Technology Officer of Intel Architecture, on directions for Internet computing; Ron Smith of the Wireless Communications and Computing Group with a major announcement on Intel's entry into the wireless communications arena; Mark Christensen of the Network Communications Group on network infrastructure and new developments from Intel, and Will Swope of the Solutions Enabling Group on software solutions for e-Business and the media-rich e-Home, and the tools and services that can make you more successful.

IDF Fall 2000 gives you an inside look at the latest in platform technologies, communications products, and software solutions. Here are the highlights.

"Powering the net"

Intel platform technology powers the net, and IDF Fall 2000 is your opportunity to learn about the latest technologies and building blocks for platforms based on IA-64, including an announcement from Intel on the roll-out of the Itanium™ processor. Intel and other industry experts will discuss IA-64 system architectures, related technologies, and the tools and resources available to help developers deliver optimized products. Sessions include an overview of the Itanium processor microarchitecture, optimizing IA-64 device drivers, understanding complex server topologies, and Intel's server and workstation chipset strategy.

The Intel® Pentium® 4 performance desktop architecture, formerly known as "Willamette," will share the IDF platform technology spotlight. IDF Spring 2000 featured a demo of this revolutionary processor at 1.5 GHz. If you thought the performance was impressive, Intel has an even more powerful demonstration planned for IDF Fall. The conference will also provide the latest information on tools and methodologies you can use to implement next-generation IA-32 products. Practical information for designers includes an in-depth discussion of Intel Pentium 4 microarchitecture and its Level 2 cache, performance tuning, motherboard design guidelines, chipset integration, and other design considerations for next-generation systems.

"Connecting the net"

Intel is expanding on its core competencies in the design and delivery of processors and chipsets to become an important building block supplier for developers of wireless communications and computing products. A number of announcements in this important technology area will be made at the conference. Technical sessions will focus on how Intel is helping to enable the transition to packet-switched networks to support the convergence of voice and data. Intel will share its vision of wireless LAN technologies and how they fit within the 2.4-GHz band, including the Bluetooth specification for "personal area networking," wireless Ethernet (802.11b) for business LANs, advances in HomeRF technology, and a look ahead to opportunities in the 5-GHz band.

In the fast-growing area of mobile commerce, Intel will showcase the building blocks of "anytime, anywhere" communication and computing, including baseband and RF chipsets and digital signal processors.

The Intel® Internet Exchange™ Architecture (IXA) is creating important new opportunities for developers of network devices and infrastructure products. OEMs will announce some of the latest communications and networking products based on IXA building blocks, including the IXP1200 network processor, the IX-API, embedded silicon, and applications from third-party developers.

This fall's IDF sessions also provide a wealth of IXA design information, including using network processors, programming the IXP1200 network processor with the IX-API, and enabling application-aware networking.

Software Solutions

IDF has more in store for software developers than ever before, including solution-based technologies for e-Business and e-Home products. Fully 40 percent of the technical sessions at IDF Fall 2000 are devoted to important software for multiple operating systems. The sessions include hands-on labs, case studies, and examples, featuring development techniques, tools, optimization methodologies, and porting tips direct from the industry's technology leaders.

The Trusted Computing Platform Alliance will detail enhancements to client computing platforms which increase trust, while retaining owner flexibility and privacy. The sessions will help OEMs, hardware, BIOS, and software vendors plan for "Trusted Client" product development and application support.

Intel is working with the industry to enable the software services and tools needed to support the latest advances in Intel Architecture for net-centric products and applications. Sessions are organized in three major topic areas:

Developing software for the e-Home includes a case study on building an e-Home today, media sharing, creating searchable video for the Internet, scalable 3D graphics, 3D graphics performance optimizations, digital TV datacasting, and other vital topics.

Tools optimization and porting provides an in-depth look at tools and techniques available to deliver high-performance, high-quality software and Internet applications based on the Intel Architecture.

The *e-Business solutions* area provides information on developing and integrating scalable, reliable, and secure end-to-end-solutions.

Summary

IDF Fall 2000 is more solutions-focused than ever before, with enabling technologies for IA-32 and IA-64 platforms, wireless communications applications and infrastructure products, and the latest information for developing rich consumer and business-critical applications and Web sites.

You'll want to be on-hand as Intel makes an announcement of major importance for developers of wireless products and OEMs unveil their latest IXA-based products.

IDF Fall 2000 has more demos than ever before, with more than 60 percent more hours in the Demo Showcase and three times more space devoted to software.

IDF has so much technical information to offer, you will need to send a team to cover all the essential hardware and software topics you need for powering and connecting the net, and energizing your company's development program in the months ahead.

More Info

Make your plans now. For information and registration details, visit Intel's IDF Web site. It provides information on San Jose, California, accommodations, in addition to details on keynotes, technical tracks, labs, and demos.

Author Bio

David Barkai joined Intel in 1996. Before assuming responsibility for IDF Conference content, he worked in the company's Microcomputer Software Lab, where he focused on applications for Intel Architecture workstations. Prior to that, David specialized in scientific and engineering supercomputing applications and conducted pioneering work on vector processors. He holds a Ph.D. in theoretical physics and has published over 20 papers in the areas of physics, numerical methods, and computer technical applications and architectures.

Vincent Merrick joined Intel in 1984. He has worked in a number of technical and marketing positions, including design engineer, test engineer, technical marketing engineer, product marketing engineer, and OEM account manager. Vince has spent the last two years as the geography customer marketing manager for the OEM Platform Solutions Division.

Columns

From the Editor

Donna Loveland
Managing Editor
Intel Corporation

Column

The Intel® Developer Forum Conference, Fall 2000 is just a month away, and Intel's developers and colleagues are building toward it with plenty of information-rich articles.

Optimizing Performance in Virtual Private Networks takes a look at VPNs, a key component of enterprise, extranet, and wide-area network security. A recent trend has been to design dedicated platforms for achieving optimal performance in VPNs. When developed on an Embedded Intel® Architecture platform, VPN cost, performance, time-to-market, flexibility, and scalability are improved.

Zero to OS in 20.2 Seconds focuses on the Intel® Rapid BIOS boot and how it improves user experience. Intel has optimized BIOS code to eliminate redundancies and streamline the boot process, reducing Power-On Self Test time by up to 25 seconds and reducing the time to Windows* launch by approximately 50 percent.

Motherboards: An Answer to Industry Demand examines the OEM's demand for a motherboard that bridges the gap between commercial motherboards and the unique requirements of embedded systems. RadiSys Corp. took this demand as a challenge and developed the industry's first Intel® Architecture-based motherboards designed with embedded systems in mind.

Serial ATA: An Evolutionary Transition delivers an update on the software-transparent "drop-in" replacement for parallel ATA that maintains compatibility with existing operating systems and drivers, adding performance headroom for years to come.

*Intel Enables Universal Plug and Play for Linux** introduces the Intel® Universal Plug and Play Software Development Kit V1.0 for Linux*, which provides a cost-effective way to integrate UPnP features in Linux-based products like Internet appliances, residential gateways, and home networking products.

The Intel® Itanium™ Processor and Security tells how several features of the Intel® Itanium™ processor, based on the IA-64 architecture, can help to speed up public-key cryptography, which is widely used in Internet-based security technologies. Its performance advancements can help IT managers eliminate the trade-off between performance and security.

Bring Your Team to Fall IDF, our cover story, offers a preview of the industry's "must-attend" event for hardware and software developers who need to keep up-to-speed on the latest Intel® Architecture platform products, wireless technologies, and software enabling tools. With the equivalent of 40 days of content packed into three Conference days, it will take a team to cover it all.

This Much We Owe to Our Friends, this month's Inside Looking In column, gives an insider's perspective on the Conference from IDF content development manager David Barkai.

As IDF approaches, *Intel Developer Update* keeps you close to the action.

Enjoy.

Author Bio

Donna Loveland is the editor of *Intel Developer Update* magazine. She joined Intel's Platform Marketing group in 1999 as the editor of Platform Solutions News. Donna began her career with Intel in 1982 as a technical editor in an advanced microprocessor development group. Since then, she's held technical and marketing positions in leading-edge technology areas ranging from stereoscopic display to digital broadcast to scalable online content. Donna has a B.A. degree in English from the University of Rochester and an M.A. in Expository Writing from the University of Iowa.

This Much We Owe to Our Friends

David Barkai
Content Development Manager
Intel Architecture Group
Intel Corporation

Column

This was the morning when I dreaded opening my in-box, dreaded finding e-mails I was reluctant to face. Last night I had the unpleasant task of informing quite a few of our friends across the industry that their proposals for a session at our conference had been declined. How do you face the engineer from that small start-up who pulled off another late night to meet our deadline, only to be turned down? How do you face the virtual acquaintance from the previous event, the one who submitted the same interesting and innovative idea again—with the same disappointing outcome?

The form message I'd sent—not having allowed for the resources for more personal, specific responses—was polite but curt. "Thank you for the time and effort. There are more submissions than we can accommodate. This is not a reflection on the quality of the work. Please don't be discouraged from considering a proposal for the next event." It gave the recipients a glimpse into a number of reasons that may have been the cause for being turned away. But, in general, they did not receive the detailed communication or attention they deserved.

There are reasons for this omission. To understand why they're not excuses you need to know something about Intel culture. In addition to always having plenty of work to do, a "project" such as reviewing external proposals and abstracts, one that's intense but limited in time, tends to be assigned without making allowances for one's "regular" job. This phenomenon results in us being very economical with the time spent on reviews and contacts. Maybe even stingy.

It's not that we want to discourage engineers from the industry from submitting their proposals to our conference—quite the contrary. The Intel Developer Forum (IDF) Conference, the event we're talking about, is an Intel-sponsored event. We'd be less than honest if we didn't declare that it's meant to serve Intel's agenda. But we'd be foolish if we didn't recognize its symbiotic qualities. At each IDF Conference we strive to increase the industry's participation as active contributors to the speakers' program. Offering a larger choice of topics and a variety of speakers adds value and popularity to the event, enhancing the attendees' experience and increasing attendance. Broad attendance, in turn, is beneficial to the industry presenter or exhibitor: IDF offers a unique stage for companies to reach their customers and the media, to mingle, and, frankly, to observe what their competition is doing. Everyone gains.

And if everyone benefits the interactions should be easy, pleasant, and smooth. Intel puts out a Call for Papers, our industry colleagues and customers respond in time to meet the deadline, a timely and thorough review takes place, and an answer is conveyed to the submitters. In theory, this is the process.

In practice, things are different. Not everyone who should be invited to submit sees the invitation. Of those who do, most send their proposals during the 12 hours before and 24 hours after the deadline. Once we see the abstracts we scramble to find the right people to review them. As a result we don't get all the reviewers' recommendations by the promised date. Few reviewers make contact with submitters to get additional information. Finally, in order to keep from dragging out our responses any further, we resort to a curt form letter.

And a bigger problem, albeit a nice one, looms in the not very distant future. The number of proposals submitted to IDF has increased by about 50 percent each time over the two most recent conferences. If submissions continue at this rate, in just one year the number of industry proposals will equal all those we generate from within Intel.

So how is it that we've managed to get away with all this and still run the most successful developers' conference in the industry? Because everyone's pretty much in the same situation. Those last-minute contributors with the great proposals are tugging for scarce support themselves, and they're working in moments stolen from their personal lives and their day jobs. In fact, quite a few of them wouldn't have time to do more than skim over a well-considered and personally crafted e-mail response from me.

We at Intel want to do better, to be more thorough and communicate with our industry colleagues on a more personal basis. The mutual benefits are undeniable. And, after all, a strong component of Intel culture is “role modeling,” living the behaviors you want others to emulate.

At the same time, we understand that establishing and maintaining open and frank exchanges between the industry and Intel is a challenge. Business considerations, such as product plans and intellectual property, make both sides guarded.

The Intel Developer Forum is an institution that brings all our relationships with the world around us into a sharp focus. It fulfills its purpose not only in great presentations and sessions but also in the casual networking and personal discussions among the attendees in the halls and session rooms. The companies and individuals who help us create a better conference for developers deserve our full and timely attention. We owe our friends—and ourselves—that much.

And, personally, I feel better when I receive fewer messages of disappointment and more like this one, from a content future IDF presenter: “Thank you very much for your response. I look forward to being contacted by the Intel track owner, and doing my best to provide an excellent presentation for the IDF.”

Author Bio

David Barkai joined Intel in 1996. Before assuming responsibility for IDF Conference content, he worked in the company’s Microcomputer Software Lab, where he focused on applications for Intel Architecture workstations. Prior to that, David specialized in scientific and engineering supercomputing applications and conducted pioneering work on vector processors. He holds a Ph.D. in theoretical physics and has published over 20 papers in the areas of physics, numerical methods, and computer technical applications and architectures.

Departments

Applied Computing

Optimizing Performance in Virtual Private Networks

David Hillyard
Communications Platform Manager
Applied Computing Products Division
Intel Corporation

Overview

As communication traffic over the Internet increases, the need for greater network security also increases. In recent years, public telecommunication lines or frame relay circuits were the only security means available. Now, Virtual Private Networks (VPNs) provide a simple way to transmit information privately, securely, and reliably over the existing unsecured public telecommunication infrastructure.

A VPN is a simple way to transmit data safely across an unsecured network. A VPN creates a tunnel from sender to receiver. The tunnel can be formed from a remote user to a LAN, from LAN to LAN, and within an intranet. The connection is set up on a need basis and broken down when the transaction is complete. This is why the private network is considered virtual. Figure 1 shows the structure of a VPN.

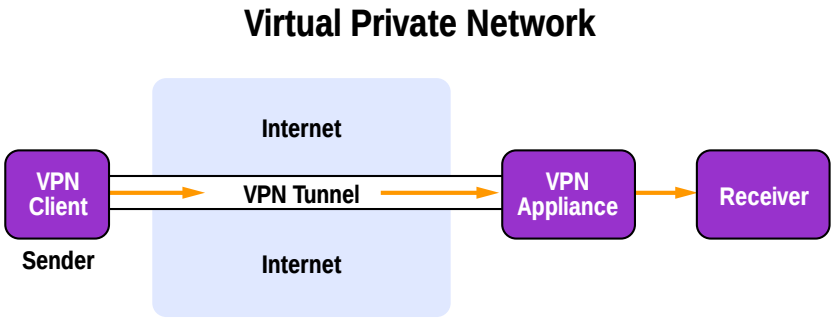

Figure 1

A VPN accomplishes the following tasks:

- Privacy—keeps unauthorized users from reading intercepted data. Privacy is accomplished by encryption and encapsulation.
- Data Integrity—guarantees the received data has not been altered.
- Access Control—has the ability to limit the amount and type of traffic entering or leaving a network.
- Non-Repudiation—enables the sender to prove to a third party that a message was sent, or the receiver to prove to a third party that the exact same message was received.
- Authentication—validates the identity of the sender.

Optimizing Network Performance

Network performance is optimally achieved when a VPN is executed on a dedicated system. On non-dedicated systems, network performance may suffer from the added responsibility of running VPN software when tasked with routing packets, performing server functions, encrypting data, and authenticating users all at the same time. Integrating VPN software on an Embedded Intel® Architecture (EIA) platform greatly increases the performance of the network by allowing all system resources to dedicate themselves to performing security tasks.

The Entry Level Reference Design

Intel provides the ideal platform to host VPN software. The Entry Level Communications Appliance Reference Design consists of a processor, chipset, Ethernet controllers, and memory. Embedded Intel Architecture (EIA) supports Microsoft Windows* 95/98/2000/NT, Linux*, Unix*, and Solaris*. The entry-level design consists of an Intel® Celeron™ processor, 440BX AGPset, and two IEEE 802.3 Fast Ethernet Ports. The two Fast Ethernet ports let the user link to a LAN and any other Internet connection device—that is, cable modem, xDSL modem, or server. The implementation described includes an RS-232 serial port for local console use. However, a VPN appliance can be administered remotely on a LAN. Figure 2 displays the entry-level design.

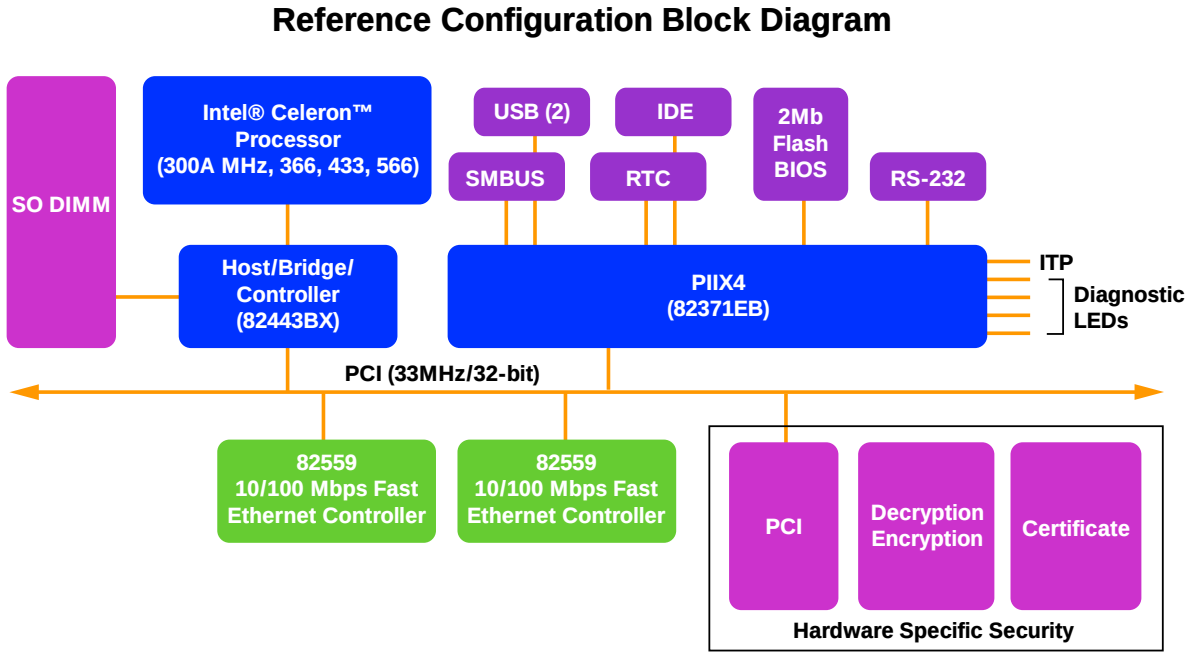

Figure 2

The sections that follow describe an Embedded Intel Architecture VPN solution that addresses the unique performance requirements of this application.

Intel® Celeron™ Processor

The processor handles encryption/decryption, IP routing, and authentication. A 300-MHz processor can handle approximately 1,500 to 2,000 tunnel connections. The Intel® Celeron™ processor implements a Dynamic Execution microarchitecture and executes MMX™ media technology instructions for enhanced media and communication performance.

The Celeron processor uses the same multi-transaction, AGTL+ system bus used in the Intel® Pentium® II processor. The Celeron processor supports multiple low-power states such as AutoHALT, Stop-Grant, Sleep, and Deep Sleep to conserve power during idle times. It is based on the P6 core and is provided in a Flip-Chip Pin Grid Array (FC-PGA) package. The Celeron processor includes an integrated 128-Kbyte Level 2 cache with a separate 16-Kbyte instruction and 16-Kbyte data Level 1 caches. The Level 2 cache is capable of caching 4 Gbytes of system memory address space. Processor utilization is important for performing encryption/decryption processes.

Intel® 440BX AGPset

The Intel® 440BX AGPset is the chipset supporting the Intel Celeron processor with a 66-MHz system bus and 100-MHz SDRAM. As Intel's second-generation AGPset with Intel® Quad Port Acceleration (QPA), the Intel 440BX AGPset improves the speed of the system bus from 66 MHz to 100 MHz, while increasing the width and depth of buffers to the system bus, Accelerated Graphics Port (AGP), SDRAM, and PCI bus. In addition, the 440BX AGPset can interface with ATA/66 HDD in UDMA mode 2. The 82443BX has the following features:

- Support for single Celeron processor configuration
- 64-bit GTL+ based host bus interface
- 32-bit host address support
- 64-bit main memory interface with optimized support for SDRAM at 100 and 66/60 MHz
- 32-bit primary PCI bus interface with integrated PCI arbiter
- Extensive data buffering between all interfaces for high throughput and concurrent operations

Intel® 82559

The 82559 is Intel's second-generation, fully integrated 10BASE-T/100BASE-TX LAN solution. The 82559 consists of both the Media Access Controller (MAC) and the physical layer (PHY) interface combined into a single component solution. The combined component is packaged in a 15 mm x 15 mm, 196-lead, thin BGA package.

The 82559 provides 32-bit PCI bus high-speed data transfer without additional glue logic. Its bus master capabilities enable the component to process high-level commands and perform multiple operations, which lower CPU utilization by off-loading communication tasks from the CPU. CPU utilization is just as important as network throughput because the CPU needs considerable processing power to execute the tunneling encryption/decryption algorithm.

The 82559 also includes an interface to a serial (4-pin) EEPROM and a parallel interface to a 128-Kbyte Flash memory. The EEPROM provides power-on initialization for hardware and software configuration parameters. The parallel port can be used as either a Flash memory interface or an ISA-like interface for a modem. Combined with a Total Cost of Ownership (TCO) controller, the 82559 can help reduce the total cost of ownership in network environments. The device includes a System Management Bus (SMB) interface enabling the TCO controller to communicate with a management agent on the network.

Operating System Software

VPN systems can be designed to use a variety of operating systems. In this reference design, Linux is selected as the operating system due to its low cost, its open configuration model, and a small system footprint. To jump-start the LAN interface in VPN appliance design, visit the Intel Developer Web site, where you can find additional technical information or a free Linux driver for the 82559 Ethernet controller.

Summary

A broader variety of communications and the growth of the Internet have sparked a concern for greater network security. Virtual Private Networks (VPNs) provide a key component of enterprise, extranet, and wide-area network security. A recent trend has been to design dedicated platforms for achieving optimal performance in VPNs. When developed on an Embedded Intel® Architecture platform, their cost, performance, time-to-market (TTM), flexibility, and scalability are improved.

More Info

You can find the following at the Developer Web site:

- Download of the Entry Level Communications Appliance Reference
- Additional information on the Intel® Celeron™ processor
- Additional information on the Intel® 440BX AGPset
- Details on Intel 82559
- A free Linux driver or related technical information for the 82559 Ethernet controller
- Intel's Software Assistant

Author Bio

David Hillyard directs communications platform strategy activities for Intel's Applied Computing Products Division. He joined Intel in 1989, and for the last several years he has been involved with emerging communications services and products, including hardware, software, and system-level programs and initiatives that support this market segment. Before joining Intel, David served in various system-level engineering capacities at Motorola, Inc. and Digital Equipment Corp.

Motherboards: An Answer to Industry Demand

Peter Mitchell
Marketing Director
Embedded Motherboard Operation
RadiSys Corporation

Overview

When there are hidden costs involved, a bargain is not much of a bargain anymore. This holds very true for the embedded systems market. As OEMs think about the total cost of ownership of their embedded systems, they must consider several things: the hardware and software costs; the cost of deploying, supporting, upgrading, and managing the equipment; and the time and money spent in lost productivity and lost opportunity when a system has to be reconfigured.

Today, commercial motherboard prices can look like bargains until you start adding in the total cost of ownership and short product life cycles. Oftentimes, OEMs end up paying more for so-called "bargain motherboards."

In the typical commercial motherboard market, OEMs must contend with several issues including high fall-out rates, limited customer/technical support after the sale, upgrades without customer notification, high internal qualification costs, large inventory purchase requirements, and the ever-recurring reconfiguration process. All those factors make motherboards some of the most expensive OEM equipment in the embedded market.

Other problems plaguing the industry include stiff regulations from governmental agencies that monitor the gaming, medical, military, and banking industries. These agencies require a whole new testing process once a motherboard reconfiguration has been made—all while your customer sits back and waits.

Additionally, the OEMs' customer needs equipment that is available for a long time. For example, conventional and digital X-ray, computed tomography, magnetic resonance, and ultrasound equipment require the same embedded system (motherboard) for the life of the product. Medical companies need products that are designed and manufactured to last for many years. Once a product is cleared by the Food and Drug Administration (FDA), companies want to market it for several years before making any changes and going before the FDA again.

Furthermore, the embedded market requires nonstandard application support from motherboard manufacturers. Many OEMs need compatibility testing to ensure they are choosing the correct solution. Lower priced products typically don't offer this kind of testing and customer support, which is required by today's OEMs.

It is easy to see why there is such a demand for a motherboard designed specifically for the embedded systems market. RadiSys Corp. took this demand as a challenge and developed the industry's first Intel® Architecture-based motherboards designed for embedded systems.

Win-Win for OEMs

Until recently, there were no motherboard products to address the specific needs of the embedded market. With the latest embedded motherboard family from RadiSys, that is on course to change. Endura* is the new family of motherboards that bridges the gap between commercial motherboards and the unique requirements of embedded systems. This motherboard family fills a huge demand in the market, since other motherboards are not designed and supported with embedded applications in mind.

Because of its long-term experience in the embedded market, RadiSys, an initial member of the Intel® Applied Computer Platform Provider Program, predicts its Endura motherboard family will signal the start of exciting market demand for embedded motherboards, in much the same way that commercial motherboards created mass markets for desktop PCs. RadiSys essentially addresses all the demands of the embedded motherboard market with a product family and value proposition that combine high quality, change notification, product life cycle management, engineering support, and extended life.

RadiSys is able to combine high-quality design and manufacturing with low-cost motherboard solutions. Design and manufacturing are intimately linked via techniques of “design for manufacture” and “design for test,” ensuring optimum efficiency in the manufacturing process. With extensive knowledge of the electronics components industry, long-term vendor relationships, and continual market benchmarking, product designs are driven toward market-linked targets. By using a high-volume, low-cost manufacturing process, RadiSys is able to offer very competitively priced embedded motherboard products.

In today's embedded market, manufacturers often confuse the term *longevity* with *short-lived* or *end-of-life (EOL).* Some motherboard manufacturers think longevity means 18 months, while others will go as far as saying two years. RadiSys knows what terms like *EOL* or *reconfiguration* mean to OEM business, and have designed the embedded motherboard family with a life span of up to five years. By firmly standing behind this goal, RadiSys is committed to making sure the motherboards OEM customers buy will be available now and in the future, will be of the highest quality, and will be supported by our customer service staff.

Extended lifetime is achieved using various methods. First, all core motherboard designs are based on a chipset from Intel’s embedded roadmap. Second, time-to-market design and manufacturing ensures that the market window for each new generation of technology is maximized. When the end-of-life of a product does eventually occur, RadiSys makes sure that customers receive formal end-of-life notification. Typically, this is a six-month notification period followed by a six-month order window. OEMs then have plenty of time to transition their products to a newer generation motherboard.

Unannounced changes in BIOS revisions or the revision of major silicon components on commercial motherboards are not uncommon. In some instances, the first time an OEM observes a change is during an inspection in manufacturing or, in the worst case, where a manufacturing or field failure occurs because of an incompatible revision change. RadiSys protects OEMs from these issues by tracking revision changes of all motherboard components and issuing a Product Change Notification to all customers. By prior arrangement, customers can participate in a Change Request Program whereby RadiSys will seek approval from the customer before introducing a change to the motherboard already in use. This is particularly important for OEMs who may have to make software updates when a major silicon component switches to a new revision. In the event of such a revision change, RadiSys will work with the customer to ensure that an appropriate solution can be found.

"We believe the introduction of RadiSys Endura motherboards is what the embedded market has been awaiting for quite some time," says Ron Dilbeck, Vice President and General Manager, RadiSys Corp. "This type of embedded motherboard product is ideal for embedded applications where the application doesn't require the high price and high availability solutions. RadiSys embedded motherboards let embedded customers get away from using the expensive CompactPCI*****, VME, and slot card architectures if they only need motherboard technology for their specific applications."

Product Overview

The Endura motherboards cover all popular form factors, offering a wide range of solutions from entry-level to performance-level computing platforms. These include ATX, micro ATX, FlexATX and micro NLX form factor motherboards that enable system designers to select an appropriate form factor to suit their needs, including high slot requirements, ISA slot support, small footprint, and low-profile chassis to a plug-in motherboard.

RadiSys motherboards are based on the wide range of Intel® Pentium® III and Celeron™ processors in the Pin Grid Array 370 (PGA370) package, and include a rich feature set of integrated technologies such as 10/100 Ethernet, a PanelLink* flat-panel controller, and integrated VGA. On some form factors, ISA slots are provided for legacy application support. This is an important feature for many OEMs who have proprietary ISA-based cards. Offering each motherboard with a number of population options providing customers with the flexibility to choose the most cost-effective solution for their application expands product options.

Although all RadiSys motherboards are based on industry-standard form factors and designs, RadiSys is able to add enhancements such as fan speed control or GPIO lines that offer extended feature sets to embedded system designers. Fan speed control is available on the most recent motherboards. During periods of inactivity, the processor generates less heat, allowing the system fans to be slowed down. This can extend the life of the system fans, improving the overall MTBF of the system. Super I/O controllers also have unused GPIO signals (originally intended for a MIDI/joystick interface). These signals are brought out to a header on the motherboard so they can be used as additional control lines

in an embedded system—a useful feature requested by many customers. Additional connectors provide the means to cable internally inside a chassis to keyboard, mouse, flat panel, and audio interfaces.

A degree of "soft customization" is also possible for high-volume customers. This may be customer-specific part number labeling or a custom BIOS that auto-configures or provides a default set of CMOS parameters. As a standard, customers are able to include their own logo in the BIOS so that it can be displayed on the startup screen.

Summary

Until now, OEMs were essentially stuck with selecting a motherboard with low quality, inferior customer support, and a shorter-than-acceptable life cycle. RadiSys Endura motherboards, based on the Intel Architecture, have changed all that. They offer OEM customers what they want: a lower total cost of ownership through longer life, proven and predictable quality, engineering change order notification, compatibility testing, customization, and the flexibility of personalized boot screens.

More Info

Visit the RadiSys Web site for detailed information on these OEM products:

- SF810 Micro NLX
- EM440 ATX
- EMB1

Author Bio

Peter Mitchell, marketing director for the Embedded Motherboard Operation, joined RadiSys in October 1999, bringing his sought-after experience in OEM business-to-business marketing. Prior to that he spent 16 years in the PC industry with the UK PC manufacturer Apricot Computers and with Mitsubishi Electric Corporation, where he held positions in both research and development, and sales and marketing. Peter holds an M.B.A. and a Ph.D. in engineering from the Universities of Wolverhampton and Cambridge.

Initiatives and Technologies

Zero to OS in 20.2 Seconds

Justin Whitney
Senior Product Marketing Engineer
Desktop Boards/System SW Marketing
Intel Corporation

Overview

The typical home PC user's first impression of a new system is the boot process. A beep, some disk activity, a splash screen, and more than 60 seconds later the typical PC is ready to go. Many users take advantage of this interval to grab a second cup of coffee or check their phone messages. "Type-A" personalities have been known to leave their PCs on all the time.

For a typical PC, approximately half of the boot interval is used by the BIOS to initialize the display and input devices, perform Power-On Self Test (POST) functions, scan system buses and main memory, and so on. In an effort to improve the user experience by shrinking the boot interval, Intel has optimized BIOS code to eliminate redundancies and streamline the boot process.

In PCs with Intel® Rapid BIOS Boot, Windows Millennium Edition* operating system, and selected hardware, boot times as short as 20 seconds have been observed (bench tests performed at OEM Platform Solutions Division, Intel Corporation, June 15, 2000. See "Test Results" section for bench test details.). It's an advance in PC ease of use with important benefits for OEMs, integrators, and end users alike.

Faster Boots

One of the ways to make PCs more "appliance-like" and user-friendly is to cut down the boot interval. Three major factors affect boot time.

BIOS optimization. Intel Rapid BIOS Boot reduces POST time by about 25 seconds, reducing the time to Windows* launch by approximately 50 percent (bench tests performed at OEM Platform Solutions Division, Intel Corporation, June 15, 2000). See "Test Results" section for bench test details.

Hardware configuration. Monitor warm-up and synchronization rates, hard drive spin-up, and CD-ROM initialization all take time. For this reason, building a fast-booting PC is also a matter of choosing the right hardware. A legacy-reduced PC with a fast spin-up hard drive is recommended. Having the floppy drive as the first boot device can add several seconds to boot time.

Operating system optimization. With Microsoft's forthcoming Windows Millennium Edition (Windows Me) operating system, the time needed to boot the OS has been cut to about 15 seconds, compared to 30 seconds for Windows 98. See "Test Results" section for bench test details.

Test Results

The bar graph in Figure 1 illustrates the results of a comparison between a PC with Intel Rapid BIOS Boot and an otherwise identical system with a generic BIOS. Both systems were equipped with Windows Millennium Edition operating system, a 733-MHz Intel® Pentium® III processor, 128 Mbytes of RAM, and a 5,400-RPM hard disk. The tests were repeated for Windows 98. Boot times may vary with system configuration. (Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel® products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance.)

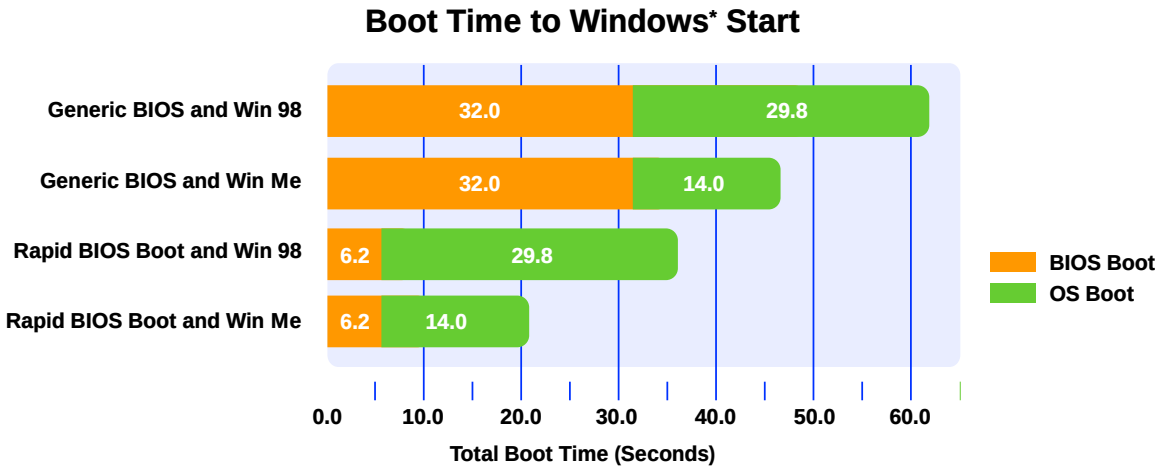

In this test, Intel® Rapid BIOS Boot resulted in an 80 percent reduction in POST time and a 56 percent reduction in the time to Windows Millenium Edition* launch

Figure 1

No Trade-offs

Improved boot performance was achieved through BIOS code optimizations not trade-offs in functionality. Here is a summary.

Improving task parallelism. In traditional versions of BIOS, the tasks of initializing the hard disk, initializing the CD-ROM, scanning the Universal Serial Bus, displaying the splash screen, and clearing system memory occur in series. Initiating these tasks to run in parallel can save time.

Removing redundant code. Legacy BIOS code can contain instructions for functions that may also be performed by the operating system during the boot process. Eliminating these redundant functions from BIOS can save more time.

Improving efficiency. Most of the time, a system boots from the hard drive rather than from the floppy drive. By default, Intel Rapid BIOS Boot does not perform a seek for floppy media thereby saving a few seconds during the Power-On Self Test.

Legacy reduction. Disabling or eliminating the floppy drive entirely can also save boot time.

Express BIOS Update

For argument's sake, let's say you have an older Intel® Desktop Board and would like to upgrade the BIOS to Rapid BIOS Boot. What if that system is legacy reduced and doesn't have a floppy drive? Intel® Express BIOS Update is a Windows-based utility available from Intel's Developer site which dramatically simplifies the task of updating system BIOS.

Intel Express BIOS Update is a 1-Mbyte executable file that can launch from within the operating system on the user's PC. The user is no longer required to boot to DOS, install the file on a floppy disk, or perform multiple system reboots. After downloading Express BIOS Update, simply double-clicking on the application leads the user through an installation wizard which automatically shuts down the system, flashes the BIOS, and reboots to the OS.

Other than its inherent ease of use, a significant advantage of Intel Express BIOS Update for system integrators and resellers is that it can simplify or reduce tech support requirements by supplying an automated OS present utility for systems with the Intel® 810, 815, and 820 chipsets. Tech support would typically have to step the user through the complicated tasks of creating a bootable floppy disk, extracting the BIOS files to a temporary directory, copying those files to the floppy, booting to DOS, running the Update program, and rebooting the system.

Additional Benefits

For OEMs. Manufacturers typically perform a number of tests that require a system reboot. Time savings in the boot process can save significant amounts of money in testing and fixture costs, or permit manufacturers to conduct additional tests and increase throughput. Using Intel Rapid BIOS Boot also enables system integrators to meet or exceed the seven-second BIOS POST standard contained in the PC 2001 System Design Guide.

For end users. Many end users and even some resellers have the perception that slow boot performance equals slow system performance, even in some of today's best-performing PCs. Therefore, reducing boot times results in increased user satisfaction.

Summary

Intel Rapid BIOS Boot can reduce POST time by more than 80 percent and reduce the time to Windows launch by more than 50 percent. See "Test Results" section for bench test details.

Boot intervals as short as 20.2 seconds have been observed in PCs equipped with Intel Rapid BIOS Boot, Windows Me, and optimized hardware.

For optimum boot performance, Intel recommends the use of Intel Desktop Boards that support Intel Rapid BIOS Boot, the optimization of system and hardware drivers, an optimized operating system, and hardware components with the fastest available spin-up times. Together, these steps will improve ease of use and make the PC experience more rewarding for end users.

More Info

For technical details on BIOS and faster booting PCs, read the BIOS Requirements section in Chapter 3 (PC System) of the PC 2001 Design Guide. In the near future, OEM Platform Solutions Division will release a white paper describing the specifics of optimizing the BIOS as it relates to boot speed.

Use Intel Express BIOS Update to download the latest BIOS files for upgrading compliant systems to Intel Rapid BIOS Boot. You can download a BIOS Upgrade at the Intel® Desktop Boards Web site.

Intel Developer Forum Conference, Fall 2000 (IDF) is in San Jose, California this August. If you are in town, stop by and see first hand how Intel Rapid BIOS Boot and Express BIOS Update have made updating and using PCs faster and easier. Visit the IDF area of the Developer Web site for more details.

Author Bio

Justin Whitney is a senior product marketing engineer within the Intel Architecture Marketing Group. His project responsibilities have included Intel Rapid BIOS Boot, Intel Express BIOS Update, and the Superconducting Liquid Cryogenic Level Sensor, for which he holds a patent. Justin holds a B.S. in mechanical engineering from Northwestern University.

Serial ATA: An Evolutionary Transition

Bill Colson
Marketing Manager
Intel Architecture Labs
Intel Corporation

Overview

The parallel ATA (AT Attachment) specification has defined the standard storage interface for PCs since the protocol was introduced in the 1980s. Parallel ATA has maintained its pre-eminence for three primary reasons: 1) low cost, 2) virtually universal operating system support, and 3) the ability of the specification to evolve to higher speed and performance while maintaining backward compatibility with older ATA devices. From its original speed of just 3 Mbytes/second, parallel ATA has moved up to a burst data transfer rate of 66 Mbytes/second. The latest generation of the interface, Ultra ATA-100, goes even further, with a burst data transfer rate of 100 Mbytes/second.

While ATA has enjoyed an illustrious track record, the specification is now showing its age. Parallel ATA imposes some serious design issues on today's developers, including a 5-volt signaling requirement, high pin count, and serious cabling headaches.

Intel Architecture Labs, in conjunction with the industry working group promoters and contributors, has developed the Serial ATA solution to overcome these design limitations while enabling the storage interface to scale with the growing media rate demands of PC platforms. Serial ATA is designed to be a software-transparent "drop-in" replacement for parallel ATA that maintains compatibility with existing operating systems and drivers, adding performance headroom for years to come. For these reasons, Serial ATA has attracted industry-wide support from PC OEMs, semiconductor manufacturers, and leading vendors of storage devices and software.

End of Parallel

All good things must come to an end, and parallel ATA is no exception. Here are the main reasons parallel ATA technology has run out of evolutionary headroom:

- *High pin count*. Parallel ATA requires 26 pins per channel. Multiplied by two channels, that's 56 pins per channel, plus power/ground.
- *High voltage*. Parallel ATA requires 5-volt transceivers, which impose integration problems with new silicon processes.
- *Cable problems*. The 80-conductor cable required to support parallel ATA is relatively expensive and unwieldy to route inside the PC chassis. Flat ribbon cable can interfere with air flow and cooling. To make matters worse, the ribbon header connector can be difficult to seat in the system board and the storage device during the assembly processes, and can lead to reliability and support issues.
- *Performance issues*. Evolving parallel ATA beyond 100 Mbytes/second could require the implementation of technical enhancements, including low voltage differential (LVD) signaling.

With approximately two years of lead time required to design and implement a new storage interface generation, this is a logical time for parallel ATA to evolve into a standard that can continue to support the price/performance requirements of desktop and mobile PC platforms. The software transparency and backward compatibility of Serial ATA specification makes it an even more compelling and timely solution for the industry.

Benefits of Serial ATA

The Serial ATA specification is designed to replace parallel ATA with a software-transparent interface for "inside the box" storage. It reduces voltage and pin count requirements and can be implemented with thin and easy to route cables.

- *Low voltage requirement*. Serial ATA requires only 500 millivolts (mV) peak-peak to support new silicon processes and higher integration.
- *Lower pin count*. Reducing the pin count helps reduce board real estate requirements and enables more reliable connections on the board and the storage device.
- *Higher integration*. Unlike parallel ATA, Serial ATA does not require 5-volt tolerant transceivers that can pose a hindrance to higher integration which utilize today's and future silicon processes.

Improved Air Flow with Serial ATA Cable

Ribbon Cable **Serial ATA Cable**

Figure 1

- *Better cabling*. As shown in Figure 1, Serial ATA replaces cumbersome ribbon cable with thin and flexible cable that can be up to one meter in length. Simpler cabling can reduce manufacturing cost, improve the reliability of cable connections, improve system cooling, and provide system designers more room inside the chassis.
- *Ideal for mobile*. The reduced power requirements make Serial ATA ideal for use in mobile PCs, enabling OEMs to use one standard interface across a broad line of cost-effective platforms.
- *Cost-effective*. At introduction, the cost of Serial ATA is expected to be competitive with today's parallel ATA implementations.
- *Software compatibility*. Serial ATA is a drop-in replacement for parallel ATA. There is no impact to the existing infrastructure of software and drivers.

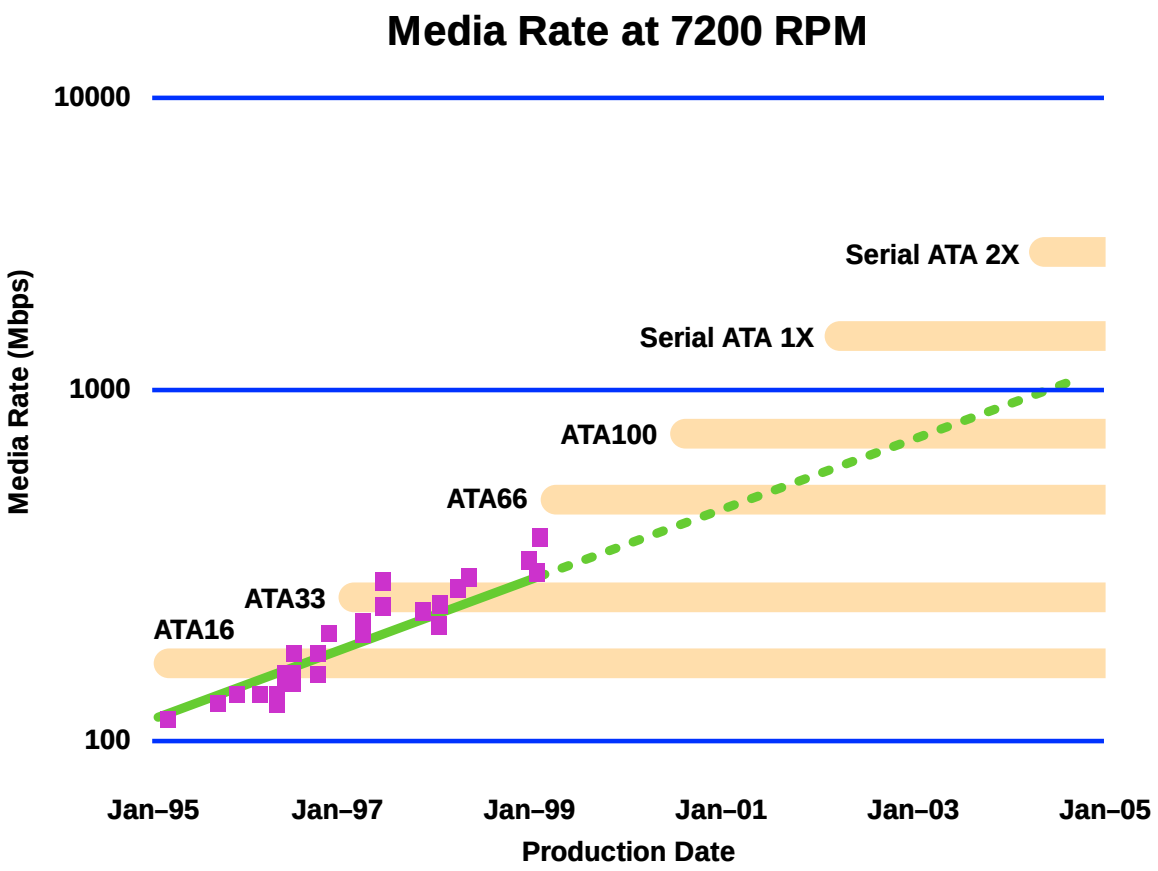

Figure 2

- *Strong roadmap*. As Figure 2 illustrates, the Serial ATA roadmap is currently projected to support scalability in successive generations of Serial ATA technology, beginning with a data rate of 150 Mbytes/second and evolving to data rates of 300 Mbytes/second and beyond.
- *Superset Features*. Serial ATA supports a superset of features including dynamic plug/unplug, addressing of drives larger than 137 Gbytes, first-party Direct Memory Access (DMA) and efficient command queuing. It's important to note that these features are not required for baseline functionality, and they will require additional software support.

When?

The 0.9 version of the Serial ATA specification is scheduled to be released at the Intel Developer Forum Conference, Fall 2000 (IDF). The goal is to have the 1.0 version of the specification completed by the end of this year. The introduction of first-generation Serial ATA devices capable of supporting 1.5-Gbit/second signaling is anticipated in mid-2001.

It is anticipated that transitional solutions will require discrete Serial ATA components in concert with chipsets that support Parallel ATA devices. Serial ATA chipsets are now under development by several third-party vendors. Attending the IDF Conference Fall 2000 may provide an opportunity to see the first Serial ATA chipsets from these suppliers.

As a transition strategy, developers may also use serial and parallel dongles to adapt parallel devices to a serial controller or adapt serial devices to a parallel controller.

Summary

The familiar parallel ATA interface has been the primary “inside the box” storage interface for more than 10 years. For a variety of reasons, including relatively high voltage requirements, cabling issues, and lack of performance headroom, a new solution is now required in order to keep pace with the growing media rate of desktop and mobile PCs.

In a development role similar to the one it took with PCI and USB technologies, Intel Architecture Labs is leading an industry working group that has developed a new interface specification designed as a drop-in replacement for the parallel ATA interface. Known as Serial ATA, this new serial interface is designed to overcome the limitations of parallel ATA while launching a roadmap that could drive 10 more years of storage interface performance evolution.

Thanks to its combination of software transparency, low cost, scalability, and design flexibility, Serial ATA has attracted widespread industry support through the Serial ATA Working Group. Membership in the working group provides developers with early access to the specification and influence over finalization and information, which can speed the transition to this innovative and cost-effective new storage interface technology.

More Info

More information about the Serial ATA specification and the benefits of membership in the Serial ATA Working Group is available on the Serial ATA Working Group Web site.

Plan to attend the Intel Developer Forum Conference Fall 2000 for the latest information on implementing the Serial ATA specification. Watch Intel’s Developer Web site for details.

Author Bio

In his role as marketing manager for the Intel Architecture Labs, Bill Colson oversees a wide range of technologies in the areas of platform architecture, voice/speech/media/data, and Internet services. During his 16-year career at Intel, Bill has been involved in the development of the Multibus I and Multibus II single board computers, and desktop and server technologies, including the first Intel® 386 PC baseboard. Bill holds a patent in the field of server management. He is a recipient of the Intel Achievement Award, and a former member of the IEEE and Next-Generation I/O Forum. A frequent contributor to *Intel Developer Update*, he has written articles for *EE Times* and *Electronic Design.* Bill holds a B.S. in computer systems engineering and electrical engineering from the Oregon Institute of Technology.

Intel Enables Universal Plug and Play for Linux*

Andrew Liu
Product Marketing Manager
Intel Architecture Labs
Intel Corporation

Steven Conner
Technical Marketing Manager
Intel Architecture Labs
Intel Corporation

Overview

Universal Plug and Play (UPnP) is an open networking architecture for peer-to-peer device discovery and control. In last month's issue we introduced UPnP and elaborated on ways UPnP offers consumer value through compelling connected experiences—utilizing easy, affordable Internet access and connectivity—anywhere in the home.

Intel's commitment to UPnP technology is strong. The Intel Architecture Labs in Oregon developed the Connected.Home initiative, which works to establish affordable, easy-to-use, and broadly deployed network connectivity to and within the home. This initiative directly drives the creation of open industry specifications required to deliver on this vision, including UPnP. Intel's e-Home and Ease-of-Use initiatives, designed to bring user-friendly device and appliance management into the home, also support enabling seamless operation of products in the e-Home of the future.

At the June Summit of the Universal Plug and Play Forum in Seattle, Intel introduced the Intel® Universal Plug and Play Software Development Kit V1.0 for Linux*. This kit (SDK) provides developers with a cost-effective way to speed time-to-market for UPnP-compliant Linux-based products such as Internet appliances, residential gateways, and home networking products.

UPnP Software Developer Kit (SDK)

The Intel® UPnP SDK provides a comprehensive Application Programming Interface (API) and Linux source code for the implementation of UPnP-compliant control points and devices. The SDK will be deployed and tested on the Linux operating system and is interoperable with UPnP implementations on Windows*. The SDK, which is compliant with version 1.0 of the UPnP specification, will include header files, commented source code, documentation, and an open source software license.

The SDK provides a complete UPnP reference implementation, and supports both controlled devices and control points. Although the SDK has been developed and tested on Linux, it is written using standard POSIX APIs and should compile on any POSIX-compliant operating system.

Key Features

The Intel UPnP SDK features a single, unified API that is used by both control points and devices. The SDK includes an XML abstraction using Document Object Model (DOM) Level 1, and an optional integrated Web server. The SDK supports both synchronous and asynchronous operations, as well as a variety of callback functions.

The SDK executes callbacks in a thread context, allowing memory allocation and asynchronous callback with other threads. The default callback is registered with the client or device, though each asynchronous function may specify a different callback function.

Figure 1 illustrates the Intel® Linux UPnP SDK, open source with no restrictions. There are three main components to effective implementation of UPnP. At the top of the stack is the UPnP Client or Service Application, provided by the customer, shown in blue. This runs on top of the SDK to implement functionality of specific services, such as gateway commands to enable or disable an Internet connection. Samples of both client and server code are included in the kit, but the application itself is left to the device or software developer to implement.

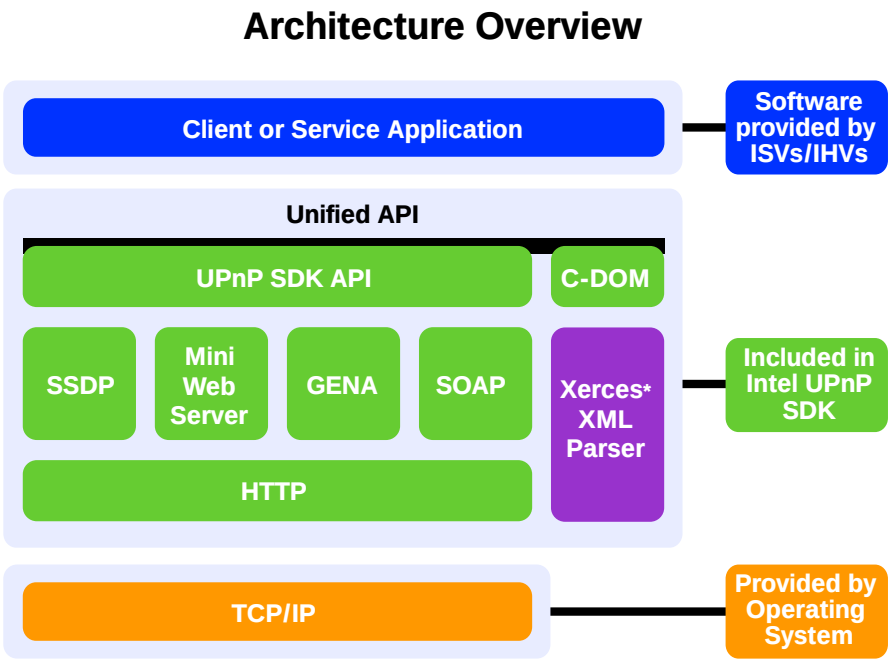

Figure 1

The next layer in Figure 1 is the SDK, featuring eight integrated components. The seven green boxes in the gray SDK layer are included in the Intel SDK; the eighth box (shown in purple) is the XML parser. For convenience, the Xerces* XML parser is included with the Intel SDK package. Although an XML parser is required by the SDK, the developer may choose to replace the Xerces parser with another DOM-compliant parser.

The Application Programming Interface (API) shown in Figure 1 abstracts the details of the core UPnP protocols away from the client or service, providing access to all functionalities provided by UPnP.

The *Unified API* provides functionality for both clients and servers. Devices (which behave as servers in UPnP) are able to advertise the services they provide and watch for queries from other devices and control points. Control points (which behave as clients in UPnP) use the API to send out searches for devices and services, watch for notification from devices and services, and control remote devices. In both instances, the API makes it possible for the application software to operate without having to directly interpret the SSDP, GENA, and SOAP. Through this single, unified API it is possible for a single device on a network to act as both a UPnP control point and controlled device.

The Simple Service Discovery Protocol (SSDP) implements the discovery phase of UPnP, allowing clients to multicast search requests for devices or services and receive the response from those searches. Likewise, SSDP enables devices to advertise their existence and services to clients when they are connected to the network.

The Mini Web Server handles standard HTTP GET requests, allowing XML description, presentation, and control documents to be served from a device or control point. This module manages the locations of documents that are available and implements the actual streaming of the data via the HTTP protocol.

The General Event Notification Architecture (GENA) implements the eventing phase of UPnP. Control points use this module to subscribe or unsubscribe to events from device services. Devices receive subscribe and unsubscribe requests from control points via this module and generate appropriate event notifications.

The Simple Object Access Protocol (SOAP) provides the control phase of UPnP. Clients use this module to generate the appropriate XML documents to retrieve or change the state tables of device services. Devices use this module to decode the control requests and generate the correct response.

The HTTP layer accepts all HTTP requests for SSDP, SOAP, GENA, and the Mini Web server over TCP and UDP connections. It will parse all HTTP headers and hand the connection over to the appropriate module for processing.

The Xerces XML parser* shown in the purple box is used extensively in the UPnP SDK. All description and control data is communicated by XML documents in UPnP, so the XML parser is used by both the core UPnP protocols and by the client or service software. Examples include GENA utilization of XML for event change notifications and SOAP use of XML to format control requests and responses. While Xerces was chosen as a convenient XML parser to include with the UPnP SDK, this parser is abstracted through an API to enable substitution with an alternate XML parser, if a developer so chooses.

The C-DOM module exports a C-style Level 1 Document Object Model interface, implemented on top of the Xerces XML parser used internally by the SDK. C-DOM provides an abstract interface to manipulate XML documents without having to deal with the XML directly or the details of how it is stored in memory. Since Xerces uses a C++ API, the C-DOM layer is included on top of the Xerces parser to give a unified C-API with the UPnP SDK.

Finally, at the bottom layer of the stack depicted in yellow, is the *TCP/IP* stack provided in the Linux OS. While this stack is also not part of the UPnP SDK, it is shown here to provide a complete illustration of core UPnP protocols. The UPnP SDK assumes the TCP/IP stack is provided by the operating system.

UPnP SDK v1.0 API

Figure 2 provides a table of C-functions that are included in the UPnP SDK API. These functions are called by a customer-supplied device or control point application to register with the SDK, declare callback functions (also provided by the customer) that will be called when queries and control requests are received over the network, and to initiate communication with remote UPnP devices and control points. Sample source code is included with the SDK to demonstrate the effective use of this API and associated callback functions.

C-functions Included in the UPnP SDK API

	Devices	Control Points	Both
General	`UpnpRegisterRootDevice, UpnpUnRegisterRootDevice`	`UpnpRegisterClient, UpnpUnRegisterClient`	`UpnpInit`
Discovery	`UpnpSendAdvertisement`	`UpnpSearchAsync`	
Eventing	`UpnpAcceptSubscription, UpnpNotify`	`UpnpRenewSubscription, UpnpRenewSubscriptionAsync, UpnpSubscribe, UpnpSubscribeAsync, UpnpUnSubscribe, UpnpUnSubscribeAsync`	
Control		`UpnpGetServiceVarStatus, UpnpGetServiceVarStatusAsync, UpnpSendAction, UpnpSendActionAsync`	

Figure 2.

Learn More About UPnP

To encourage quick adoption of the UPnP SDK, Intel will present demonstrations of this exciting new technology at two upcoming events.

Intel Developer Forum Conference, Fall 2000 (IDF) scheduled for August 22 through 24, 2000 in San Jose, California, will feature a UPnP Implementers Workshop led by the engineers who created the Linux SDK. The workshop will focus on applying the code to real applications. A UPnP Pavilion at IDF will provide attendees the opportunities to see the latest demos from early implementers. You can also get the latest updates on the multi-platform, multi-vendor interoperability and development tools. To learn more about IDF and to register see the More Info section of this article.

The UPnP Interoperability Plugfest, held at Intel Oregon, is scheduled for September 20 through 21, 2000. This will be a hands-on engineering event focused on multi-vendor interoperability. Developers are invited to bring their latest devices and control points, and be ready to rebuild and test code on site. To register for the Plugfest see the More Info section of this article.

Summary

The Intel Universal Plug and Play Software Development Kit V1.0 for Linux enables developer support of UPnP and could dramatically simplify use of Linux-based Web appliances and e-Home devices. Equally important to standards-based technology and implementation of ease-of-use products, the UPnP SDK is interoperable with Windows* operating system software.

The SDK, with the Application Programming Interface and Linux source code, supports developers interested in tapping the growing e-Home industry. The SDK will be available beginning summer 2000 and is an ideal tool to support developers interested in adding UPnP capabilities to products.

More Info

To learn more about the Intel Universal Plug and Play (UPnP) Software Developer Kit (SDK) for Linux, and to receive notification of source code availability, see the UPnP area of the Intel Architecture Labs Web site.

You can get full information about the Intel Developer Forum Conference, Fall 2000 and register to attend on the Intel Developer Web site.

For the UPnP Plug Fest at Intel, visit the UPnP Web site.

Background information on Universal Plug and Play can be found in “Universal Plug and Play Connects the Home” in the July issue of the *Intel Developer Update* magazine.

For more information on the Connected.Home initiative, visit the Intel Architecture Labs Web site. Information on the Intel Ease-of-Use Initiative is available in the Technology area of the Intel Web site.

For additional information on the Document Object Model, see the W3C* Web site.

Author Bio

Andrew Liu is a product marketing manager in the Connected Consumer Initiative in the Intel Architecture Labs (IAL) and is responsible for Home Networking Infrastructure technologies, including residential gateways, Universal Plug and Play, and wireless technologies. He has been with Intel since 1993. Prior to his current position, he worked with Intel’s Anywhere in the Home, scalable platforms, and Internet media technologies. He has presented at the Intel Architecture Conference, ISPCON, the Universal Plug and Play Forum, and WinHEC. Andrew holds a B.S. degree in Electrical Engineering and Computer Science from the University of California at Berkeley, and a Master of Electrical Engineering from Cornell University.

Steven Conner is a technical marketing manager in the Connected Consumer Initiative in the Intel Architecture Labs. Steven has been with Intel since 1999 and is responsible for Home Networking Infrastructure technologies, focusing on Residential Gateways and Universal Plug and Play. In his previous work, he focused on Linux software development. Before joining Intel, he worked in the area of computer vision research and development. Steven holds a B.S. degree in Electrical Engineering and a M.S. degree in Electrical and Computer Engineering from the University of Arizona.

Executive Summary

The tremendous growth of the Internet has resulted in an explosion of e-Commerce Web sites. As e-Commerce continues to grow, security becomes a major issue for IT managers. To enhance security, Web sites use software that relies on security algorithms and protocols to authenticate identities, protect data, and monitor transactions.

One type of security algorithm is public-key cryptography, which is widely used in Internet-based security technologies. Common protocols such as Secure Sockets Layer (SSL), Public Key Infrastructures (PKI), and Virtual Private Networks (VPN) are examples of security technologies that rely on algorithms based on public-key cryptography. Despite the widespread reliance on these algorithms, they have a significant drawback. They are very compute-intensive and can have a significant impact on server performance. This is especially true in the case of short transactions, which are typical of e-Commerce.

The Intel® Itanium™ processor, based on the IA-64 architecture, has several features that can help to speed up public-key cryptography. Parallel instruction issue, multiple execution units, 64-bit Integer Multiply-Add instruction, and a large register-set enable high performance on public-key cryptography algorithms. Using optimistic performance estimates for competitive products (based on theoretical analyses of their respective instruction sets and architecture) and comparing with actual performance on the Itanium processor, the Itanium processor is nearly 10 times faster than the Sun UltraSPARC* III processor. (See the full Intel white paper "Intel® Itanium™ Processor: High Performance On Security Algorithms" in the IA-64 architecture area of the Developer Web site). Table 1 compares decryption performance between Itanium and UltraSPARC III processors.

RSA 1,024-bit Decryption Performance Comparison

Processor	Clock Frequency	RSA Decryptions Per Second
Itanium™ Processor	660 MHz	1,000*
Sun UltraSPARC* III	600 MHz	130**

Table 1.
*Measured performance
**Estimate of best-case performance based on theoretical analysis

The Intel Itanium processor provides performance leadership in executing the public-key algorithms needed for secure e-Commerce transactions, and it does so without special hardware assistance. This is an advantage because hardware solutions are costly and increase the form factor of the total server solution.

Public-key cryptography, also known as asymmetric-key cryptography, differs from symmetric-key cryptography. In symmetric-key cryptography, there is one key, and it is used by the sender for encryption and by the receiver for decryption. This approach introduces potential problems with key distribution.

In public-key (or asymmetric-key) cryptography, there are two keys: a public key known to all, and a private key known only to the key owner. This approach eliminates any key distribution problems, since the public key does not need to be hidden. The downside is that compared with symmetric-key cryptography, public-key cryptography is very compute-intensive, which inhibits its use for encrypting large volumes of data.

For this reason, symmetric-key and public-key cryptography are generally used in tandem. The keys for symmetric encryption are exchanged using public-key cryptography, and the actual message is exchanged using the symmetric keys. The most commonly used implementations of public-key encryption are the RSA public-key algorithm, patented by RSA Security Inc., and the protocol known as Secure Sockets Layer (SSL). Some other security technologies that rely on public-key cryptography are Public Key Infrastructures (PKI) and Virtual Private Networks (VPN).

Named after its inventors (Rivest, Shamir, Adelman), the RSA algorithm is the most widely used public-key algorithm. It is also computationally intensive. At the core of the computation is the exponentiation of large numbers to large powers. The encryption of a single small message may cost millions of native processor operations. There are alternatives to using the RSA algorithm for public-key encryption, but virtually all of them use a large number multiplication at their core.

To understand large number multiplication, consider computation of the quantity M^e mod n, where each of the values is a very large integer, typically 512 or 1,024 bits in length. The mod (or modulo) operation computes the remainder after division and is commonly referred to as a reduction, since it reduces the number to be less than n. Since the exponent is so large, it is obviously not feasible to complete all of the exponentiation before doing the reduction, because the resulting number would have an astronomical number of bits. In practice, the multiplication steps are interleaved with reductions after each multiplication. The long division used for reduction is very expensive if done by repeated trial division.

In 1985, a researcher named Peter Montgomery published a method for doing the multiplication and reduction that does not require trial division. Now known as the Montgomery Product, this method allows the multiplication and reduction to be done as a series of multiplications and additions. The Montgomery Product is used widely for doing modular exponentiation. This calculation takes up about half the time required to complete a short secure transaction typical of e-Commerce.

The Itanium processor has several features that help to accelerate the performance of the Montgomery Product. These include parallel instruction issue, multiple execution units (four integer, four multimedia, two floating-point, two memory, and three branch units), large register set, and the 64-bit Integer Multiply-Add Instruction (xma).

The 64-bit Integer Multiply-Add instruction produces the upper or lower 64-bit product, executes on either floating-point execution unit, and is fully pipelined. The large register set allows the Montgomery Product operands and intermediate results to be held in registers rather than in memory.

Because of the potential security vulnerabilities posed by e-Commerce, Web sites need the protection of security-specific software, even though such software is widely known to degrade server performance. Until recently, IT managers have been forced to live with the performance trade-off or resort to expensive hardware solutions. Now, performance advancements provided by the Itanium processor enable the integration of additional security features into application software without compromising server performance. This means IT managers can finally eliminate the trade-off between performance and security.

Since public-key algorithms are the core of many security applications, Itanium processors should significantly improve the performance of these applications and enable the integration of additional security features into application software without compromising server performance. This goes a long way toward eliminating the historical trade-offs that IT managers have been forced to make between security and performance.

More Info

For a more detailed discussion of the impact of the Intel Itanium processor on security algorithm performance, including a performance comparison with other platforms, see the Intel white paper “Intel® Itanium™ Processor: High Performance On Security Algorithms” in the IA-64 architecture area of the Developer Web site.

—End of Intel Developer Update Magazine Issue 11—